



The AI SOC analyst that investigates every alert and acts on it.

Ornis Guard is an on-premises AI SOC agent for threat detection and incident management. It triages alerts from various cybersecurity technologies around the clock, explains every verdict and contains threats.



AI

From alert overload to explainable decisions

Security teams evaluate hundreds of alerts every day. Most are false positives, yet each one consumes analyst time and lengthens dwell time for the threats that matter. Ornis Guard turns raw alert records into explainable, auditable, action-ready security cases.

Where SOC and MDR teams lose time

- **Alert fatigue:**
High volumes and repeated false positives push real threats down the queue.
- **Manual enrichment:**
Threat intelligence, asset and user context are gathered by hand, alert by alert.
- **Inconsistent triage:**
Quality depends on who is on shift and how much time they have.
- **Slow handoff:**
The gap between investigation and response is where attackers gain ground.

What changes with ornis guard

- **Full coverage:**
Every alert is triaged and investigated by specialized AI roles, 24/7.
- **Consistent verdicts:**
Each case carries a confidence score, an evidence chain and a rationale.
- **Faster response:**
Policy-driven containment shortens both MTTD and MTTR.
- **Compounding knowledge:**
Analyst and customer feedback become organizational memory

The TDIR workflow — six steps for every alert

Each stage is executed by a specialized AI role and logged in the case timeline.

1. Collect

✓ ok

Receives new alerts, filters them and eliminates duplicates.

2. Enrich

✓ 2.0s

Correlates IOCs with threat intelligence, assets and organizational context.

3. Investigate

✓ ok

Evaluates the alert step by step through specialized AI roles.

4. Validate

✓ ok

An independent quality layer audits the result before anything leaves the system.

5. Respond

● approval

Delivers the verdict via ticket, e-mail or Telegram; executes or proposes containment; closes the alert at the source.

6. Learn

✓ stored

Adds analyst and customer feedback to organizational memory.

Coverage

Every alert

Triaged and investigated — not sampled, not queued.

Explainability

Every verdict

Confidence score, evidence chain, MITRE ATT&CK mapping.

Response

Every action

Policy-gated, logged, reversible where the platform allows.

One workflow for SOC and MDR operations:

Beyond alert summaries, Ornis Guard brings analysis, threat hunting, quality control, case management, customer communication and organizational learning together in a single, auditable workflow.

An explainable decision for every alert

Each alert passes through five stages performed by specialized AI roles. Monitor the review live — see how the decision is produced and how confident the system is at every step.

Web Attack: PHP-CGI Remote Code Execution Attempt (CVE-2024-4577)

Ornis Guard ID: OG-1140191

Alert ID: 309682

IsNet

High (8/10)

Analyzed

Summary

Analyses

Timeline

Analyst comments

Investigate

Hunter Report Reviewer

COMPLETED

Aug 24, 03:55 PM

• True Positive

confidence 0.94

✓ File creation traces
no findings

✓ Persistence check
no findings

✓ Lateral movement scan
no findings

✓ C2 communication analysis
203.0.113.47 monitored

✓ No persistence detected — impact contained, evidence attached.

Five stages, five specialized AI roles

- 1

Initial assessment · Triage analyst
Classifies and prioritizes the alert; identifies affected assets and users.
- 2

Investigation plan · Investigator
Runs controlled queries in the SIEM and connected platforms, within the permitted scope.
- 3

Threat hunting · Threat hunter
Runs controlled queries in the SIEM and connected platforms, within the permitted scope.
- 4

Quality review · Quality reviewer
An independent QA role validates the evidence, the reasoning and the confidence score.
- 5

Decision and action · Responder
Delivers the reasoned verdict to the relevant channels and hands the case to the response policy.

Analysis outputs

- Verdict**
True positive or false positive, stated clearly.
- Confidence and severity**
Confidence score and severity level for prioritization.
- Rationale**
The evidence chain supporting the decision.
- Attack indicators**
Extracted IOC and IOA data.
- Framework mapping**
MITRE ATT&CK techniques and Cyber Kill Chain phase.
- Recommended action**
Classification, priority and the next response step.

Each role's reasoning is recorded in the case timeline. Analysts, auditors and customers see the same evidence chain — nothing is a black box.

Beyond SIEM: detection and response across EDR, XDR and NDR

Ornis Guard is not limited to SIEM alerts. It ingests detections and telemetry from endpoint, extended and network detection platforms, correlates them into a single case, and can run the complete MDR operation — detection, investigation and active response — for one organization or for many.

SIEM

Ingests alerts, offenses and log queries.

Responds by annotating and closing alerts at the source.

EDR

Ingests endpoint detections, process and file telemetry.

Responds by isolating hosts, terminating processes, quarantining files.

XDR

Ingests cross-domain, correlated incidents.

Responds by coordinating actions across endpoint, identity and cloud sensors.

NDR

Ingests network detections, lateral movement and C2 indicators.

Responds by blocking IPs and domains at the perimeter.

e-Mail security

Ingests phishing and malware detections.

Responds by quarantining messages and blocking senders.

Firewall / NGFW

Ingests connection and policy events.

Responds by pushing IP and domain block rules.

Built for MDR delivery

Multi-tenant organizations and domains · Per-customer policies, exceptions and knowledge · Customer communication via tickets and e-mail · Shift analysis and executive reporting · SLA-aware escalation and timeouts

Hybrid Response Mode

A decision engine between detection and action

Every verdict passes through a policy-driven decision engine. It weighs confidence, severity, asset criticality and the reversibility of the proposed action, then routes the case to one of three lanes.

Confidence score

Severity

Asset criticality

Action reversibility

Organization policy



DECISION

LANE 1 · AUTONOMOUS

Low-risk, reversible actions execute immediately.

- Block a known-malicious IP or domain
- Quarantine a phishing e-mail
- Open, enrich and close tickets

LANE 2 · APPROVAL-GATED

High-impact actions wait for oneclick analyst approval.

- Isolate an endpoint
- Terminate a process, quarantine a file
- Escalate to the customer

LANE 3 · RECOMMEND ONLY

Low confidence or outside policy: the analyst decides.

- Ambiguous or conflicting evidence
- Critical production assets
- Actions not covered by a playbook

• Approval required

REMOTE ACTION

Endpoint isolation

10.34.18.22 (FILESRV-02) will be isolated from the network via the EDR integration. Active sessions are preserved; rollback is available.

Requested by: Ornis Guard · Scope: single endpoint · Audit: enabled

Cancel

Approve

GUARDRAILS ON EVERY ACTION

- Scoped to organization and domain
- Logged with actor, reason and evidence
- Reversible where the platform supports it
- Time-boxed with automatic retry
- Confirmed back to the source system

Automated investigation, from suspicion to evidence

Threat hunting

Generate hunt queries with AI; launch, monitor, complete and log them.

```
index=auth user="mark.loward" earliest=-24h  
| stats count by action
```

- SIEM-specific query language and behavior
- Organization- and domain-based authorization scope
- Searches only permitted log sources
- Start, monitor, retrieve results and cancel
- Persistent query and result history

Threat intelligence

Extracts IP, domain, URL and hash indicators from alerts and combines provider results with organizational context.

203.0.113.47 · risk 92

cdn-update.example · risk 61

sha256 9f3a...c2 · clean

- Global and organization-specific API key pools
- Provider prioritization and health monitoring
- Normalized risk score
- Organization-specific exception rules
- Asset inventory and log source context

Threat map

Indicators, assets, users, evidence and MITRE techniques in a single relationship view.



OG-1140002

Resolved

AI: TP

309680 – Exploit: Apache Struts RCE Payload Detected on DMZ Server [T1190]

- Ticket created
2:13:00 PM · Analysis Workflow
- Notification e-mail sent Mail sent
2:13:25 PM · Analysis Workflow
- SIEM alert closed · source IP blocked
2:13:29 PM · Integration Outbox
- Resolved: known incident
2:30:37 PM · Ornis Guard

Communication and Memory

Tickets and customer feedback

Create tickets automatically from the analysis. Manage status transitions, comments, event history, external actions, escalation and timeouts. Generate knowledge records from feedback automatically.

Organizational knowledge base (RAG)

Manage records by source type, quality, validity period and alert context. Every new case draws on previous experience.




AI chat assistant

Escalate this to an L2 analyst.

L2 analysis complete: verdict upheld — True Positive.
Recommended: isolate 10.34.18.22, blacklist 203.0.113.47.
Completed in 3.4s — no human handoff needed

Threat intelligence, alert analysis, asset lookup, SIEM queries, IOC history, shift analysis — in natural language.

Manage SOC and MDR operations from a single screen DASHBOARD



Representative interface — no real customer data. Real-time statistics · Status and time charts · Shift analysis
Technical and executive reports · TR / EN / NL interface

Alert management

Search and filter; start or cancel individual or bulk reviews. Ignore, classify and escalate alerts; approve pending response actions.

Evidence and quality

View the timeline, evidence and enrichment data. Manage analyst notes, tags, quality scores and updates.

Reporting and audit

Generate CSV and PDF outputs. Track accuracy with the confusion matrix; resolve appeals; manage notifications and the semantic archive.

Security and Management

Integration management

- Connection via API key or user credentials
- Organization- and domain-based TLS validation
- Custom certificates and multi-domain technology support
- Scan interval and alert filters
- Log sources, asset inventory and bulk import

Security

- Encrypted integration credentials in the database
- Separate authentication keys between services
- Fail-secure behavior for incomplete configuration
- Only required external ports
- Hardened deployment and current TLS

Audit and continuity

- Retention policies by data class and organization
- Audit logs for all user and response actions
- System, query and scan logs
- Automatic retry of failed external actions
- Traceability for every critical operation



Keep your data within your organization

Ornis Guard — including its AI models — runs entirely within your own infrastructure. Security data never leaves the organization; every critical operation, including every response action, remains traceable and auditable.

Experience Ornis Guard in your security operations.

For a demo and more information:
ornis.com.tr
info@ornis.com · 0850 290 0 290

